

# Room for improvement in automatic image description: an error analysis

**Emiel van Miltenburg**

Vrije Universiteit Amsterdam  
emieli.van.miltenburg@vu.nl

**Desmond Elliott**

ILLC, Universiteit van Amsterdam  
d.elliott@uva.nl

## Abstract

In recent years we have seen rapid and significant progress in automatic image description but what are the open problems in this area? Most work has been evaluated using text-based similarity metrics, which only indicate that there have been improvements, without explaining what has improved. In this paper, we present a detailed error analysis of the descriptions generated by a state-of-the-art attention-based model. Our analysis operates on two levels: first we check the descriptions for accuracy, and then we categorize the types of errors we observe in the inaccurate descriptions. We find only 20% of the descriptions are free from errors, and surprisingly that 26% are unrelated to the image. Finally, we manually correct the most frequently occurring error types (e.g. gender identification) to estimate the performance reward for addressing these errors, observing gains of 0.2–1 BLEU point per type.

## 1 Introduction

Automatic image description is the task of describing an image in natural language (Bernardi et al., 2016). Recent advances in this area have been evaluated with text-based similarity metrics, such as BLEU (Papineni et al., 2002) or Meteor (Denkowski and Lavie, 2014). These metrics make it easy for researchers to benchmark the effect of their modeling decisions, but they are not informative about the strengths and weaknesses of a proposed model. This is especially true for n-gram based metrics, such as BLEU, which measure grammatical fluency and not semantic adequacy (Reiter and Belz, 2009).

In this paper, we present a coarse- and fine-grained analysis of the descriptions generated by a state-of-the-art attention-based model (Xu et al., 2015), trained on the Flickr30K dataset (Young et al., 2014). The goal of this paper is to assess the qualities of a state-of-the-art model to illustrate the recent progress in this area and the challenges ahead. The coarse analysis quantifies whether the descriptions are accurate or inaccurate, while the fine-grained analysis quantifies the types of errors observed in the descriptions. We define accurate to mean that the description is congruent with the image, without it necessarily being the “best” or most complete description. We find that 80% of the descriptions contain at least one type of inaccuracy, and that 26% are completely wrong. In addition to categorizing the errors, we perform a manual error correction study to estimate the reward for addressing these errors. We find that fixing the five most frequently occurring errors contributes between 0.2 – 1 BLEU points improvement over the baseline, for each type of error. We hope that our findings will encourage future research to address the specific errors we have observed.<sup>1</sup>

## 2 Related work

Early work on image description was evaluated with text-based similarity measures *and* a human judgment study (Bernardi et al., 2016). This type of judgment study involves asking humans to rate whether the descriptions accurately describe the image, are grammatically correct, are relevant for the image, are human-like, *inter-alia*, using a Likert-scale survey. The main criticisms of human judgment studies is they are expensive to perform and difficult to replicate without access to the same subject pool and control samples (e.g. Papineni

<sup>1</sup>All our code, data, and annotation guidelines will be available upon publication.



Figure 1: Examples of images with 1–4 errors. The annotated errors are marked in boldface.

et al. 2002; Hodosh and Hockenmaier 2016). Nevertheless, these studies are the clearest indication of the differences between models. Our coarse-grained analysis is a binarized version of the correctness scale from (Mitchell et al., 2012).

In this paper, our main focus is to provide a detailed analysis of the quality of descriptions generated by a state-of-the-art model. Our work is most closely related to Hodosh and Hockenmaier (2016), who propose an evaluation of image description systems using binary forced-choice tasks, where systems have to choose the best description for a given image. For each image, the system can choose between the original description or a manipulated description. By controlling the manipulations, the authors are able to check for weaknesses in image description systems. Their error categories partially overlap with ours, though we provide a more fine-grained typology.

### 3 Error categories

We developed a non-exhaustive categorisation of errors by inspecting the descriptions generated by an attention-based image description model (Xu et al., 2015). We trained the model on the Flickr30K dataset (Young et al., 2014), with 300D word embeddings, a 1000D GRU hidden layer (Cho et al., 2014), and ‘CONV<sub>5,4</sub>’ image features from the VGG-19 CNN (Simonyan and Zisserman, 2015). We generated 1,014 descriptions with a beam width of five hypotheses, recording a Meteor score of 17.4 on the Flickr30K test set.

In total, we identified 20 common types of errors, which we grouped into four main categories: PEOPLE, SUBJECT, OBJECT, and GENERAL. We developed annotation guidelines with examples for each type of error. Due to space constraints, we provide the annotation guidelines in the supplementary material. The error categories and types

of errors are described below.

**People** Image description models often make mistakes that are specific to the description of people. Types of errors in this category are AGE (e.g. *woman* instead of *girl*), GENDER (*man* instead of *woman*), TYPE OF CLOTHING (*shirt* instead of *jacket*), and COLOR OF CLOTHING (*red shirt* instead of *blue shirt*).

**Subject** Mistakes relating to the subject of the description. This category contains the following types of errors: WRONG when the wrong entity in the image is chosen as the subject, SIMILAR when the model mis-identifies the subject for something visually similar (e.g. *guitar* instead of *violin*), NON-EXISTENT when nothing close to the mentioned entity is present in the image, and EXTRA SUBJECT when an additional (nonexistent) entity is described along with the correct entity.

**Object** Similar to **Subject**.

**General** Mistakes that are not specific to people. Error types in this category are: STANCE for posture-related mistakes, ACTIVITY for wrongly identified activities, POSITION for mistakes in spatial relations within the image, NUMBER for counting errors (too few/many entities mentioned), SCENE/EVENT/LOCATION for mis-identifications of the scene, event, or location, COLOR for non-clothing entities that are mistakenly attributed with a color, OTHER for any unforeseen mistakes, and GENERALLY UNRELATED for descriptions that do not seem to have any relation with the image. In these cases, it is impossible for annotators to assign any error category to the description. E.g. if Figure 1a were to be described as *A dog runs through the snow*.

### 4 Annotation tasks

We define two error annotation tasks: The **coarse-grained annotation** task is a binary categorization

problem, where an annotator determines for every description whether it is accurate. The **fine-grained annotation** task is a multiclass categorization problem, given the error types presented in the previous section. Each inaccurate description is annotated with one or more error types. We can think of this task as a means to assess the *semantic edit distance* between a generated description and the closest accurate alternative.

In total, one annotator categorized all 1,014 generated descriptions into the coarse-grained groups: accurate and inaccurate descriptions. The same annotator then performed the fine-grained annotation. We validated the annotation scheme by double-annotating a random selection of 100 descriptions (10% of the data) to determine whether the annotation guidelines provide a reliable basis for annotating the errors.

#### 4.1 Results for the coarse-grained task

In the coarse-grained annotation task, 812 out of 1014 descriptions (80%) were judged to be inaccurate. We achieved a good inter-annotator agreement of Cohen’s  $\kappa=0.67$ , with an accuracy of 91%. The discrepancy between these numbers is explained by the label distribution: the INACCURATE category is so dominant that any disagreement yields a high penalty in  $\kappa$ . Out of the 100 double-annotated descriptions, the first and second annotator judged 86 and 81 descriptions to be inaccurate, with agreement on 79 descriptions.

#### 4.2 Evaluating the fine-grained annotations

In the fine-grained annotation task, we double-annotated the 79 descriptions that both annotators agreed contained at least one inaccuracy. Tables 1 and 2 show the number of errors per image, and the distribution of error types across the dataset. In total, we found 1,265 errors in 812 descriptions, which is an average of 1.56 errors / description.

Surprisingly, the most common error category is GENERALLY UNRELATED (264 times). Errors from the GENERAL and PEOPLE categories are much more frequent than the other two. Taken together, the SUBJECT category is least common. Our intuition is that this is because mistakes in decoding the subject from the language model affect the entire sentence; the choice of subject influences the probability of all subsequent words, leading to a generally unrelated sentence.

The fine-grained annotation task is inherently ambiguous because inaccurate descriptions might

| Errors    | 1   | 2   | 3  | 4  |
|-----------|-----|-----|----|----|
| Frequency | 486 | 221 | 83 | 22 |

Table 1: The distribution of error annotations.

be corrected in many different ways. Figure 1a illustrates this ambiguity. The generated description for this image is given in Ex. (1a). This description could either be corrected to (1b) or (1c), depending on whether one assumes the mistake is in the color or the type of clothing.

- (1) a. A woman in a **red shirt** is standing in front of a building
- b. A woman in a **black shirt** is standing ...
- c. A woman in a **red skirt** is standing ...

Subjectivity and ambiguity are inherent to the task of image description; describing an image in one simple sentence means that you have to make a choice about what to include in your description. But this subjectivity also means that it is difficult to provide a proper intrinsic evaluation for the annotation task: different choices about how to describe an image may be equally valid. To quantify the extent of this issue, we treat the double annotation for the fine-grained task as a retrieval problem, i.e. how many error types are also found by the second annotator? This experiment achieves a precision of 0.54, with a recall of 0.55. Based on this observation, we decided to carry out an *extrinsic* evaluation: how useful are the fine-grained annotations for guiding future research on model development? We discuss this evaluation below.

## 5 Correcting the errors

Now we have observed the frequency of each type of error, we can ask: would there be a positive effect if a model could address these errors? We selected the five most common error types (excluding GENERALLY UNRELATED), and manually corrected each error *without* looking at the reference descriptions. If a description is annotated with multiple errors, we only correct the relevant error. We tried to be conservative in our corrections; e.g. for COLOR OF CLOTHING errors, if the system wrote e.g. *white shirt* instead of *checkered/leopard print/... shirt*, we left the description untouched, rather than insert the pattern. For the ACTIVITY errors, we tried to change as little as possible but editing the activity often also entails changing the object as well. For example, a sentence that read

| Type                 | Count | Type                | Count | Type                 | Count |
|----------------------|-------|---------------------|-------|----------------------|-------|
| generally unrelated  | 264   | non-existent object | 47    | color                | 14    |
| color of clothing    | 195   | age                 | 40    | non-existent subject | 11    |
| activity             | 168   | stance              | 38    | wrong-object         | 7     |
| type of clothing     | 104   | position            | 37    | similar-subject      | 3     |
| gender               | 98    | extra subject       | 34    | extra object         | 1     |
| scene/event/location | 91    | similar-object      | 31    | wrong-subject        | 1     |
| number               | 61    | other               | 20    |                      |       |

Table 2: Number of times each error was annotated in our fine-grained analysis.

*A man in a suit is holding a sign.* was changed to *A man in a suit is talking.* because the man wasn't holding anything and leaving out the object would produce an ungrammatical sentence. If a change would entail completely re-ordering the sentence, we leave the generated description untouched.

Table 3 presents the BLEU and Meteor scores for the validation set before and after correction. For example, after only correcting the colors of clothing, we find a one-point improvement for the BLEU score with respect to the original model.

| Type                 | BLEU | $\Delta$ | Meteor | $\Delta$ |
|----------------------|------|----------|--------|----------|
| Baseline             | 17.8 | —        | 17.2   | —        |
| Color of clothing    | 18.8 | 1.0      | 17.5   | 0.3      |
| Activity             | 18.5 | 0.7      | 17.7   | 0.5      |
| Type of clothing     | 18.1 | 0.3      | 17.4   | 0.2      |
| Gender               | 18.6 | 0.8      | 17.6   | 0.4      |
| Scene/event/location | 18.0 | 0.2      | 17.4   | 0.2      |

Table 3: Error categories and the BLEU-4 and Meteor scores after correcting the errors.  $\Delta$  indicates improvement in the scores between the modified descriptions and the original descriptions.

We did not investigate whether these effects are cumulative, i.e. what happens if we correct *all* errors. Presumably, they are cumulative, but this task is not suitable for such an investigation because the corrections need to be restrictions in order for the improvement estimation to be accurate. If we allowed annotators to correct all the errors in a sentence, we would be giving them *carte blanche* to rewrite everything, turning the analysis into an evaluation of human performance.

## 6 Conclusion

In this paper we provided an extensive error analysis for image descriptions generated by a state-of-the-art attention-based model. Our main contributions are: (1) Providing a taxonomy of com-

mon errors in automatically generated image descriptions. (2) Quantifying the weaknesses of the model. We posit that any model with a similar architecture will have similar weaknesses. (3) Quantifying the possible improvement of this model if those weaknesses are addressed.

We focused on the nature of the inaccurate descriptions, and looked at different errors that these contain. But what about the accurate descriptions? The descriptions that *are* accurate, are also much more general than the human descriptions, which usually include small, but salient details. We propose the following rule: if the majority of the human descriptions comments on an aspect of the image that is not addressed by a generated description, then that aspect could be improved. We plan to explore the consequences of this in future work.

We see two other perspectives to build on the observations from this paper. **Automated error analysis:** As noted earlier, Hodosh and Hockenmaier (2016) carried out a study in which they evaluate image description models using binary forced-choice tasks, where models have to choose which description best describes a particular image. The choices are carefully manipulated, so that each task evaluates the model's performance in one area (e.g. recognizing scenes). Our taxonomy of errors could be used to extend the range of available tasks, for example with a task to evaluate the use of color terms; **Extending existing models:** Table 3 provides an indication of how much a model could improve by incorporating a dedicated module to detect color, actions, type of clothing, gender, and scenes. We expect that our work will encourage researchers in vision & language to investigate this possibility. More generally, we hope that our taxonomy of error types will help others to go beyond similarity-based metrics, and to look at their model's output through a qualitative lens.

## Acknowledgments

EM is supported by the Netherlands Organization for Scientific Research (NWO) via the Spinoza-prize awarded to Piek Vossen (SPI 30-673, 2014-2019). DE is supported by NWO Vici grant nr. 277-89-002 awarded to Khalil Sima'an.

## References

- Raffaella Bernardi, Ruket Cakici, Desmond Elliott, Aykut Erdem, Erkut Erdem, Nazli Ikizler-Cinbis, Frank Keller, Adrian Muscat, and Barbara Plank. 2016. Automatic description generation from images: A survey of models, datasets, and evaluation measures. *Journal of Artificial Intelligence Research* 55:409–442.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*. pages 1724–1734.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*.
- Micah Hodosh and Julia Hockenmaier. 2016. Focused evaluation for image description with binary forced-choice tasks. In *Workshop on Vision and Language, Annual Meeting of the Association for Computational Linguistics*. volume 3.
- Margaret Mitchell, Xufeng Han, Jesse Dodge, Alyssa Mensch, Amit Goyal, Alex Berg, Kota Yamaguchi, Tamara Berg, Karl Stratos, and Hal Daumé III. 2012. Midge: Generating image descriptions from computer vision detections. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, pages 747–756.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, PA, USA, ACL '02, pages 311–318. <https://doi.org/10.3115/1073083.1073135>.
- Ehud Reiter and Anja Belz. 2009. An investigation into the validity of some metrics for automatically evaluating natural language generation systems. *Computational Linguistics* 35(4):529–558.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of The 32nd International Conference on Machine Learning*. pages 2048–2057.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* 2:67–78.

## A Annotation Guidelines

### A.1 Introduction

This document provides guidelines for the annotation of automatically generated image descriptions. Our goal is to assess the semantic competence of image description models. In other words: are the descriptions at least ‘technically’ correct? This is a low bar, as we ignore fluency and usefulness, which are also desirable properties for an NLG system. We define two tasks:

1. **A binary decision task**, where annotators judge whether or not a description is congruent with an image.
2. **A categorization task**, where annotators select error categories that apply for incongruent descriptions.

These tasks are strongly related: if a description is incongruent, it should fall into one of the error categories, and vice versa. Hence, annotators for either task need to be familiar with our taxonomy of errors.

### A.2 Error categories

All our error categories are provided in Table 4. There are four main categories: People, Subject, Object, and General. I tried to strike a balance between specificity and amount of categories. No doubt some of these could be further subcategorized, but more categories means the annotation task might become overwhelming.

#### A.2.1 Short description

Here’s a short description of each category, and each of the subcategories. The next subsection provides examples for each of these.

**People** Image description models often make mistakes that are specific to the description of people. Subcategories are AGE (e.g. *woman* instead of



| People            | Subject       | Object       | General  | General              |
|-------------------|---------------|--------------|----------|----------------------|
| Age               | Wrong         | Wrong        | Stance   | Scene/event/location |
| Gender            | Similar       | Similar      | Activity | Other                |
| Type of clothing  | Inexistent    | Inexistent   | Position | Color                |
| Color of clothing | Extra subject | Extra object | Number   | Generally unrelated  |

Table 4: Error categories for incongruent image descriptions. The organization of these categories corresponds to the organization of the categories in the annotation environment.

*girl*), GENDER (*man* instead of *woman*), TYPE OF CLOTHING (*shirt* instead of *jacket*), and COLOR OF CLOTHING (*red shirt* instead of *blue shirt*).

**Subject** Mistakes relating to the subject of the description. We use the following subcategories: WRONG when the wrong entity in the image is chosen as the subject, SIMILAR when the image description system mis-identifies the subject for something visually similar (e.g. *guitar* instead of *violin*), INEXISTENT when nothing close to the mentioned entity is present in the image, and EXTRA SUBJECT/OBJECT when an additional (nonexistent) entity is mentioned besides the correct entity.

**Object** See **subject**.

**General** Mistakes that are not specific to people. The subcategories are as follows: STANCE for posture-related mistakes, ACTIVITY for wrongly identified activities, POSITION for mistakes in spatial relations within the image, NUMBER for any counting errors (too few/many entities mentioned), SCENE/EVENT/LOCATION for misidentifications of the scene, event, or location, COLOR for non-clothing entities that are mistakenly said to have a particular color, OTHER for any unforeseen mistakes, and GENERALLY UNRELATED for generally unrelated descriptions, that are beyond repair. This is usually the case when more than 2–3 error (sub)categories are applicable.

### A.2.2 Examples



A **man** is climbing a rock  
Category: Age



A **girl** playing soccer  
Category: Gender



A girl in a **yellow shirt** is standing on the beach  
Category: Type of clothing



A man in a **blue shirt** and blue jeans is working on a ladder  
Category: Color of clothing




A **boy** jumps over a hurdle  
Category: Wrong subject



A **woman in a blue shirt** is standing in front of a blue car

Category: Inexistent subject



A young girl in a white shirt is playing with a **guitar**

Category: Inexistent object



**Two police officers** are posing for a picture

Category: Similar subject, number



A man with a tennis racket **and a tennis racket**

Category: Extra object



A man in a white shirt **and a man in a white shirt** are preparing food

Category: Extra subject




A man in a brown jacket is **standing** in front of a wall

Category: Stance




A young boy is holding a **little girl**

Category: Wrong object



A black dog **runs through** the grass

Category: Activity



A man is playing a **guitar**

Category: Similar object



**Two men** are playing instruments

Category: Number





A little girl in a white dress is walking **in** the water

Category: Position



A man in a white shirt and a woman in a white shirt are standing **in a hallway**

Category: Scene/event/location



A black **and white** dog is playing in the snow

Category: Color



A group of people standing in the snow

Category: Generally unrelated



A group of people are standing in **a fire**

Category: Other

### A.2.3 Important contrasts

While the categories are fairly straightforward, there are cases where it is easy to get confused between a pair of categories. Here are additional guidelines for difficult cases that I have encountered.

- **STANCE versus ACTIVITY:** Use the former when the difference is static, e.g. *standing* vs. *sitting*. Use the latter if the difference is dynamic, e.g. *standing* versus *walking*.
- **SCENE/EVENT/LOCATION versus POSITION:** Use the former when the surroundings are not correct. Use the latter when position within the surroundings is not correct.
- **EXTRA SUBJECT/OBJECT versus NUMBER:** Use the former when the subject/object is wrongfully extended with a conjunction (e.g. *and a woman in a white shirt*). Use the latter when there's a general mismatch in number (*a, one, two, three, a group of*).
- **SIMILAR OBJECT versus POSITION:** This conflict arises in cases where e.g. *... is sitting on a bench* is used instead of *... is sitting on a chair*. In all these cases, use *similar object*. (Even if there is an actual bench in the image.)

### A.3 Task descriptions & instructions

Now that we have seen the different error categories, we can describe the two main tasks as follows:

**Task 1: Congruency** Judge whether the generated description is congruent (no error categories apply) or incongruent (at least one error category applies).

**Task 2: Categorizing incongruent descriptions** Annotate the 'semantic edit distance' between the generated description and the



closest valid description that you can imagine. Tick all the error categories corresponding to the things you would have to change. If the generated description is unrelated to the image, or if you feel that there are too many changes necessary to get to a valid description, select GENERALLY UNRELATED.

The threshold for when a description is generally unrelated is undefined. In general, I feel like type/color of clothing don't really hurt the relation between description and image as much as e.g. having the wrong verb. So it all comes down to your intuition.

#### **A.4 Evaluation: correcting the errors**

This is a separate task that serves both as an evaluation of Task 2, and as an indication of system performance if all errors identified in Task 2 are addressed. The correction task works as follows.

1. Select an error type to correct. E.g. COLOR OF CLOTHING.
2. Go through all images annotated with this type, and correct *only* the relevant error.
3. When all relevant errors are corrected, we evaluate the results using BLEU/Meteor.

It is important for this task to be conservative in editing the descriptions. Try to change as little as possible. If a change would require restructuring the entire sentence, leave the description as it is. We'd rather underestimate than overestimate the improvement from fixing the errors. Otherwise we'd just be evaluating how good humans are at writing descriptions. So e.g. for colors, *only* change color terms into other color terms. For gender, only change *man* ↔ *woman* and *boy* ↔ *girl*, not *man* ↔ *girl*. That would be changing the age along with the gender.